

PR: More than Meets the Eye

Anderson Rocha Siome Goldenstein
Instituto de Computação
Universidade Estadual de Campinas
CEP 13084-851, Campinas, SP – Brasil
{anderson.rocha, siome}@ic.unicamp.br

Abstract

In this paper, we introduce a new image descriptor for broad Image Categorization, the Progressive Randomization (PR), that uses perturbations on the values of the Least Significant Bits (LSB) of images. We show that different classes of images have a distinct behavior under our methodology, and that using statistical descriptors of LSB occurrences and enough training examples, the method already performs as well or better than comparable existing techniques in the literature. With few training examples, PR still has good separability, and its accuracy increases with the size of the training set. We validate our method using four image databases with different categories.

1. Introduction

Image Categorization is the body of techniques that distinguish between image classes, pointing out the global semantic type of an image. Here, we want to distinguish the class of an image (e.g., *Indoors* from *Outdoors*), or the type of an object in restricted domains (e.g., vegetables in a supermarket cashier). A typical scenario for a consumer application is to group a photo album, automatically, according to classes or to automate a supermarket cashier.

Recently, there has been a lot of activity in the area of *Image Categorization*. Previous approaches have considered patterns in color, edge and texture properties to differentiate photographs of real scenes from photographs of art [3]; low- and middle-level features integrated by a Bayesian network to distinguish indoor from outdoor images [11, 18]; histogram and DCT coefficients features to differentiate city images from landscape images [19]; bag of features [13, 10, 8]; and first- and higher-order wavelet statistics to distinguish photographs from photo-realistic images [12]. In addition, Fei-Fei et al. [4] have used a Bayesian approach to unsupervised one-shot learning of object categories; Heidemann [9] has presented an

approach to establish image categories using histograms, colors and shape descriptors with an unsupervised learning method. Oliva and Torralba [15] have proposed a computational model for scene recognition using perceptual dimensions, coined Spatial Envelope, such as naturalness, openness, roughness, expansion and ruggedness. Bosch et al [2]. have presented an unsupervised scene recognition procedure using probabilistic Latent Semantic Analysis (pLSA). Vogel and Schiele [20] have presented a semantic typicality measure for natural scene categorization.

The Progressive Randomization (PR) image descriptor captures statistical properties of the images' LSB channel, information that is invisible to the naked eye. The PR descriptor has three stages: the randomization process, that progressively perturbs the LSB value of a selected number of pixels; the selection of feature regions, that makes global descriptors work locally; and the statistical descriptors analysis, that finds a set of measurements to describe the image.

We show that different classes of images have a distinct behavior under PR descriptor. With enough training examples, the method already performs as well or better than comparable existing techniques in the literature. With fewer training examples, the method also provide very good performance and interesting properties for association with other image descriptors for broad images categorization purposes such as those described earlier in this paper.

In order to validate the multiclass classification as a complete self-contained classification procedure, we have used a 40,000-image database with 12,000 outdoors, 10,000 indoors, 13,500 art photographs, and 4,500 computer generated images (CGIs). with three different classification approaches: Naïve Bayes Dense Codes using the binary classifier Bagging of Linear Discriminant Analysis (NBDC-BLDA), All Pairs majority voting of the binary classifier BLDA (All-Pairs-BLDA), and SVMs [1]. In addition, we have tested the PR descriptor in three other categorization scenarios: one to provide another interpretation of the first experiment, one for 3,354 FreeFoto images categorization

into nine classes and finally, one for categorization of 2,950 image of fruits into 15 classes.

2. Progressive Randomization descriptor (PR)

Small perturbations in the LSB channel are imperceptible to humans [21] but are statistically detectable for image analysis. Some researchers have used LSBs for Steganalysis [6, 17]. Here, we introduce the Progressive Randomization descriptor for Image Categorization. It is a new image descriptor that captures the difference between broad-image classes using the statistical artifacts inserted during the perturbation process. Our experiments in Section 3 demonstrate that different types of images do have a distinct behavior under the PR.

Algorithm 1 summarizes the three stages of PR applied to Image Categorization: the randomization process (Section 2.2); the selection of feature regions (Section 2.3); and the statistical descriptors analysis (Section 2.4).

Algorithm 1 The PR descriptor

Require: Input image I ; Percentages $P = \{P_1, \dots, P_n\}$;

- 1: **Randomization:** perform n LSB pixel disturbances of the original image ▷ Sec. 2.2

$$\{O_i\}_{i=0\dots n} = \{I, T(I, P_1), \dots, T(I, P_n)\}.$$

- 2: **Region selection:** select r feature regions of each image $i \in \{O_i\}_{i=0\dots n}$ ▷ Sec. 2.3

$$\{O_{ij}\}_{\substack{i=0\dots n, \\ j=1\dots r}} = \{O_{01}, \dots, O_{nr}\}.$$

- 3: **Statistical descriptors:** calculate m descriptors for each region ▷ Sec. 2.4

$$\{d_{ijk}\} = \{d_k(O_{ij})\}_{\substack{i=0\dots n, \\ j=1\dots r, \\ k=1\dots m}}.$$

- 4: Use $\{d_{ijk}\} \in \mathbb{R}^{(n+1) \times r \times m}$ in your favorite machine learning black box.
-

2.1. Pixel perturbation

Let \mathbf{x} be a Bernoulli distributed random variable with $Prob\{\mathbf{x} = 0\} = Prob\{\mathbf{x} = 1\} = 1/2$, B be a sequence of bits composed by independent trials of \mathbf{x} , p be a percentage, and S be a random set of pixels of an input image.

Given an input image I of $|I|$ pixels, we define the LSB pixel perturbation $T(I, p)$ the process of substitution of the LSBs of S of size $p \times |I|$ according to the bit sequence B . Consider a pixel $px_i \in S$ and an associated bit $b_i \in B$

$$\mathcal{L}(px_i) \leftarrow b_i \text{ for all } px_i \in S. \quad (1)$$

where $\mathcal{L}(px_i)$ is the LSB of the pixel px_i .

2.2. The randomization process

Given an original image I as input, the randomization process consists in the progressive application $I, T(I, P_1), \dots, T(I, P_n)$ of LSB pixel disturbances. The process returns n images that only differ in the LSB from the original image, and are identical to the naked eye.

The $T(I, P_i)$ transformations are perturbations of different percentages of the available LSBs. Here, we use $n = 6$ where $P = \{1\%, 5\%, 10\%, 25\%, 50\%, 75\%\}$, $P_i \in P$ denotes the relative sizes of the set of selected pixels S . The greater the LSB pixel disturbance, the greater the resulting LSB entropy of the transformation (Figure 1).




Figure 1. Average entropy behavior for images of each category under the PR descriptor.

2.3. Feature region selection

Local image properties do not show up under a global analysis [21]. We use statistical descriptors on local regions to capture the changing dynamics of the statistical artifacts inserted during the randomization process (Section 2.2).

Given an image I , we use r regions with size $l \times l$ pixels to produce localized statistical descriptors. In Figure 2, we show the $m = 8$ overlapping regions we use in this paper. The curse of dimensionality keeps us from adding too many regions — we have found out, experimentally, that eight regions are a good tradeoff.



Figure 2. The eight overlapping regions used in the experiments.

2.4. Statistical descriptors

The LSB perturbation procedure changes the contents of a selected number of pixels and induces local changes of pixel statistics. An L -bit pixel spans 2^L possible values, and has 2^{L-1} classes of invariance under pixel perturbations (Section 2.1). Let’s call these invariant classes *pair of values* (PoV).

When we disturb all the available LSBs in S with a sequence B , the distribution of 0/1 values of a PoV will be the same as in B . The statistical analysis compares the theoretically expected frequency distribution of the PoVs with the observed ones after the perturbation process. We apply the χ^2 (chi-squared test) [22] and U_T (Ueli Maurer Universal Test) [14] to analyze the images.

2.4.1 χ^2 test

The χ^2 test [5] compares two histograms f^{obs} and f^{exp} . Histogram f^{obs} represents the observations and f^{exp} represents the expected histogram. The procedure computes the sum of the square differences of f^{obs} and f^{exp} divided by f^{exp} ,

$$\chi^2 = \sum_i \frac{(f_i^{obs} - f_i^{exp})^2}{f_i^{exp}}. \quad (2)$$

2.4.2 Ueli test

The Ueli test (U_T) [14] is an effective way to evaluate the randomness of a given sequence of numbers. U_T splits an input data S into n blocks. For each block b_i , it analyzes each of the $n-1$ remaining blocks, looks for the most recent occurrence of b_i , and takes the log of the summed temporal occurrences. Let $B(S) = (b_1, b_2, \dots, b_N)$ be a set of n blocks such that $\cup_{i} b_i = S$. Let $|b_i| = L$ be the block size for each i and $|B(S)| = N$ be the number of blocks. We define $U_T : B(S) \rightarrow \mathbb{R}^+$ as a function

$$U_T(B(S)) = \frac{1}{K} \sum_{i=Q}^{Q+K} \ln A(b_i), \quad (3)$$

where K is the number of analyzed bits (e.g., $K = N$), Q is a shift in $B(S)$ (e.g., $Q = \frac{K}{10}$ [14]), and

$$A(b_i) = \begin{cases} i & \exists i' \in \mathbb{N}, i' < i \rightarrow b_{i'} = b_i, \\ \min\{i' : b_{i'} = b_i\} & \text{otherwise.} \end{cases}$$

3. Experiments and results

In this section, we describe how we train, test and validate PR descriptor for *Image Categorization*. We validate the multiclass classification as a complete self-contained classification procedure in **Experiment 1**. In that experiment, we use a 40,000-image database with

12,000 outdoors, 10,000 indoors, 13,500 art photographs, and 4,500 computer generated images (CGIs) with three different classification approaches.

The images in **Experiment 1** come from five main sources: Mark Harden’s Artchive¹, the European Web Gallery of Art², FreeFoto³, Berkeley CalPhotos⁴, and from The Internet Ray Tracing Competition (IRTC)⁵. Figure 3 show some examples of each category.



Figure 3. Outdoors \times Indoors \times CGIs \times Arts.

We also validate the PR descriptor in three other categorization scenarios. In **Experiment 2**, we provide another interpretation of **Experiment 1** using a second carefully assembled image database. In **Experiment 3**, we perform a 9-class image categorization using 3,354 FreeFoto photographs. Finally, **Experiment 4**, we perform a 15-class image categorization using 2,950 images of fruits.

3.1. Experiment 1

We compare our method to the state of the art two-class separation approaches in the literature [3, 12, 11, 18, 16] using a simple *Bagging ensemble* of Linear Discriminant Analysis (BLDA) [7]. Furthermore, we also perform multiclass image-categorization, separating *Outdoor photographs*, *Art images*, *Photorealistic Computer Generated Images*, and *Indoors photographs*. Here, we use the BLDA classifier with 13 iterations and 10-fold cross validation.

3.1.1 Two-class classification

Cutzu et al. [3] have addressed the problem of differentiating photographs of real scenes from photographs of art works. They validated over a database with 6,000 photographs from FreeFoto and 6,000 photographs from Mark Harden’s Artchive and from Indiana Image Collection⁶.

The authors have used color and intensity edges, color variation, saturation, and Gabor features in a complex classifier. We use a similar image set reported in [3]. We have selected 12,000 photographs and 13,500 art photographs totaling 25,500 images.

¹<http://www.artchive.com>

²<http://www.wga.hu>

³<http://www.freefoto.com>

⁴<http://calphotos.berkeley.edu>

⁵<http://www.irtc.org>

⁶<http://www.dlib.indiana.edu/collections/dido>

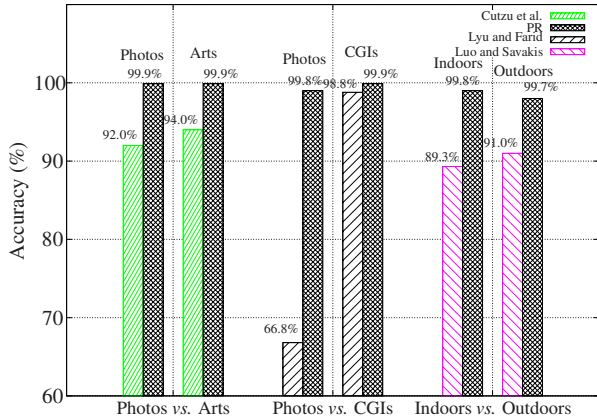


Figure 4. **Experiment 1.** PR descriptor used to binary *Image Categorization* using 10-fold cross-validation.

Lyu and Farid [12] have used a statistical model based on first- and higher-order wavelet statistics to reveal significant differences of photographs and photorealistic images. They have used photographs from FreeFoto and photorealistic images from IRTC and from Raph 3D Artists.

We use almost the same image set reported in [12]. Therefore, we have used only images from FreeFoto, IRTC and Raph sources, 7,500 photographs and 4,700 photorealistic images, totalizing 12,200.

Luo and Savakis [11, 18] have associated texture and color information about sky and grass to differentiate indoors and outdoors images. They have used a Kodak image database not freely available. Payne and Singh [16] have used edge informations to differentiate indoors from outdoors images in a personal image collection.

PR distinguishes *Photographs* from *Art* images with an average accuracy of $\frac{\mu_1 + \mu_2}{2} = 99.9\%$, *photographs* from *CGI* images with an average accuracy of $\frac{\mu_1 + \mu_2}{2} = 99.9\%$ and *Indoors* from *Outdoors* images with an average accuracy of $\frac{\mu_1 + \mu_2}{2} = 99.7\%$.

3.1.2 Multiclass classification

The PR approach creates a single descriptor that works for different separation tasks. It is suitable for multiclass broad image categorization such as the four classes *Indoors*, *Outdoors*, *CGIs*, and *Arts*.

In order to validate the multiclass classification, we have used three different approaches [1] that are combinations of binary classifiers: Naïve Bayes Dense Codes using the binary classifier BLDA (NBDC-BLDA); All Pairs majority voting of the binary classifier BLDA (All-Pairs-BLDA); and Support Vector Machines (SVMs). LibSVM uses an internal mechanism that put together all 1×1 combinations of the classes and performs a majority voting in

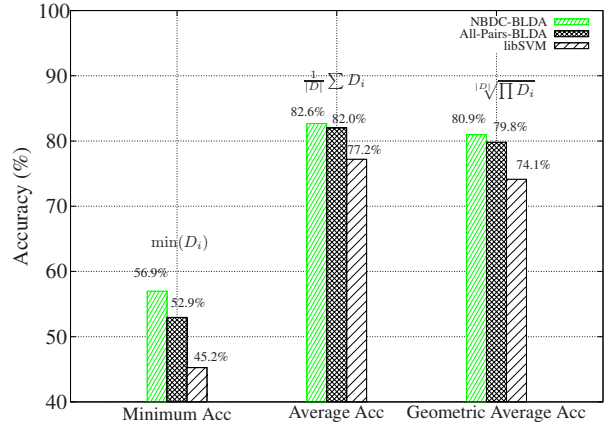


Figure 5. **Experiment 1.** Multiclass overall accuracy. D is the diagonal of the confusion matrix.

the final stage. We have used the radial basis function SVM. All-Pairs-BLDA and NBDC-BLDA use sets of binary classifiers; 1×1 combinations for All-Pairs-BLDA and $(1 \times 1, 1 \times 2, 2 \times 2, \dots)$ for NBDC-BLDA. NBDC-BLDA uses a naïve Bayes approach in which conditional probability tables of partial results come into place to improve the results. Note that, any other binary classifier could be used in All-Pairs-BLDA and NBDC-BLDA.

Tables 1-3 show the resulting classification using NBDC-BLDA, All-Pairs-BLDA, and SVMs. The diagonal represents the classification accuracy. For instance, using NBDC-BLDA multiclass approach 88.3%, of the images that represent an *Art* scenario are correctly classified, while only 6.22% of them are misclassified as *Indoors*.

The PR descriptor is independent of the multiclass approach. Figure 5 depicts the minimum accuracy of the three approaches as well as the average accuracy and the geometric average accuracy.

3.2. Experiment 2

Experiment 1 provides good performance in two-class and multi-class categorization. However, some can argue that the number of training examples is too large and that it might have suffered from bias due to different compression applied to each category, given that they come from different sources. It is important to notice that almost all images come from well known image repositories and most of them are built up from user contributions.

So, we have created a second scenario for multi-class categorization of *Indoors*, *Outdoors*, *CGIs* and *Art photographs*. In this experiment, we have manually selected 500 images for each class totalizing 2,000 images. Each category contains images from at least 75 different internet sources⁷ and there are no more than seven images from the

⁷The source path of each image can be supplied under requirement.

NBDC-BLDA Predictions				
	Arts	CGIs	Indoors	Outdoors
Arts	88.3% \pm 1.13%	5.48% \pm 0.53%	6.22% \pm 0.81%	0.00% \pm 0.00%
CGIs	30.0% \pm 2.08%	57.0% \pm 2.07%	13.0% \pm 1.22%	0.00% \pm 0.00%
Indoors	8.17% \pm 0.59%	6.34% \pm 0.37%	85.47% \pm 0.65%	0.02% \pm 0.03%
Outdoors	0.00% \pm 0.00%	0.10% \pm 0.14%	0.00% \pm 0.00%	99.9% \pm 0.14%

Table 1. **Experiment 1.** PR multiclass *Image Categorization* using NBDC-BLDA.

All-Pairs-BLDA Predictions				
	Arts	CGIs	Indoors	Outdoors
Arts	89.4% \pm 1.04%	4.41% \pm 0.49%	6.16% \pm 0.80%	0.00% \pm 0.00%
CGIs	33.66% \pm 2.36%	53.3% \pm 2.09%	13.0% \pm 1.22%	0.00% \pm 0.00%
Indoors	8.97% \pm 0.54%	5.58% \pm 0.34%	85.44% \pm 0.62%	0.01% \pm 0.03%
Outdoors	0.00% \pm 0.00%	0.06% \pm 0.07%	0.02% \pm 0.05%	99.9% \pm 0.11%

Table 2. **Experiment 1.** PR multiclass *Image Categorization* using All-Pairs-BLDA.

SVM Predictions				
	Arts	CGIs	Indoors	Outdoors
Arts	86.4% \pm 1.16%	2.10% \pm 0.43%	11.5% \pm 0.85%	0.00% \pm 0.00%
CGIs	36.2% \pm 2.32%	45.3% \pm 1.55%	18.2% \pm 1.34%	0.20% \pm 0.20%
Indoors	17.5% \pm 1.28%	4.90% \pm 0.41%	77.6% \pm 1.47%	0.00% \pm 0.00%
Outdoors	0.02% \pm 0.03%	0.37% \pm 0.18%	0.01% \pm 0.03%	99.6% \pm 0.21%

Table 3. **Experiment 1.** PR multiclass *Image Categorization* using SVM.

same place. There is no intersection among the images in this scenario with the scenario presented in **Experiment 1**. In this experiment, there are at least 400 different cameras with many different compression scenarios. Figure 6 presents the results using the All-Pairs BLDA multiclass approach with 13 iterations.

We show that the results using PR descriptor are not biased due to possible different compression levels. The results are better than the priors of each class (about 25% per class), so we show the PR descriptor provides separation among classes. Even with few training examples, the descriptor still presents a good performance. The more examples we provide in the training phase, the better the classification performance (Figure 6).

3.3. Experiment 3

Here, we select 3,354 images from FreeFoto and divide them into nine classes⁸. Figure 7 shows some examples of each category. *Sky and Clouds* category represents sunny and clear days. *Cumulonimbus Clouds* comprises images associated with heavy precipitation and thunderstorms. The other categories are self explanatory.

We do not pre-process any image. All images come from FreeFoto and were originally stored in JPEG format with 72 DPIs using similar compression levels. Figure 8 presents the results for this experiment using the All-Pairs BLDA multiclass approach with 13 iterations.

The PR descriptor generalizes from the priors (about $\frac{1}{9}$



Figure 7. 9 FreeFoto categories.

for each category). The accuracy increases with the number of training examples (left plot of Figure 8). The more images in the training phase, the more accurate the classification. This suggests that PR descriptor can be combined with other image descriptors for categorization purposes. The average standard deviation σ is below 5%. For all classes, the classification results are far above the expected priors (2σ minimum).

3.4. Experiment 4

In this experiment, we perform categorization of images of fruits and we want to show that the PR results are not biased due to camera properties. Here, we have used the same camera and setup in the capture. The JPEG compression level is the same for all images.

We personally acquired the 2,950 images at our local fruits and vegetables distribution center (CEASA), using a Canon PowerShot P1 camera, at a resolution of 1024×768 against a white background. Figure 9 depicts the 15 dif-

⁸The source path of each image can be supplied under requirement.




Figure 6. **Experiment 2.** Out × Ind × CGI × Arts using All-Pairs BLDA(13) ∴ 4 classes. *Left plot:* average performance for variable training sizes. *Right plot:* class' performance for 450-sized training sets.

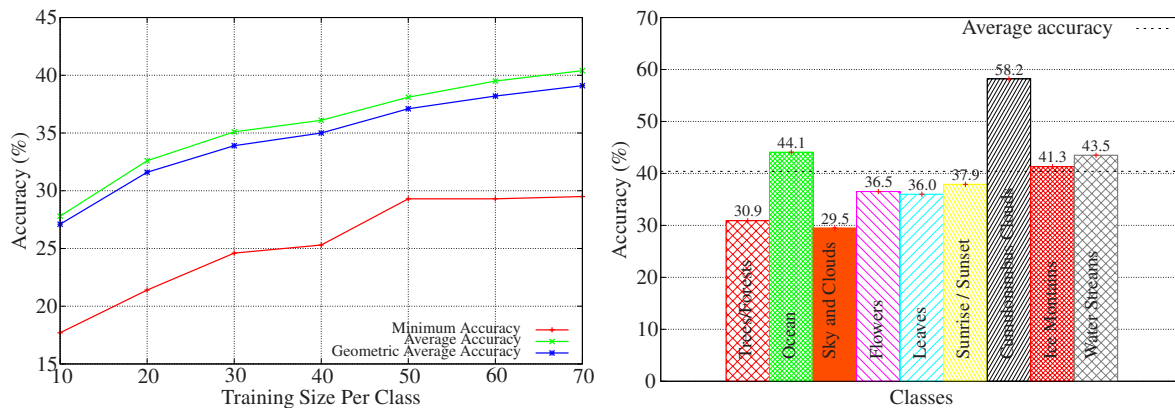


Figure 8. **Experiment 3.** FreeFoto categorization using All-Pairs BLDA(13) ∴ 9 classes. *Left plot:* average performance for variable training sizes. *Right plot:* class' performance for 70-sized training sets.

ferent classes. Even in the same category, there are many illumination differences (Figure 10).




Figure 9. 15 categories of fruits.



Figure 10. Illumination differences in the *Orange* class.

Figure 11 presents the results for this experiment using the All-Pairs BLDA multiclass approach with 13 iterations.

Clearly, PR generalizes from the priors (about $\frac{1}{15}$ for each category) and the accuracy increases with the number of training examples (left plot of Figure 11).

4. Why does PR work?

The experiments in Section 3 have demonstrated that PR descriptor can be used in broad-class image categorization. Our conjecture is that the distinct class behavior comes from the interaction between different light spectrum and the sensors during the acquisition process. That supports that fact that *Outdoors* is easier to differentiate from the




Figure 11. **Experiment 4.** Fruits categorization using All-Pairs BLDA(13) ∴ 15 classes. *Left plot:* average performance for variable training sizes. *Right plot:* class’ performance for 140-sized training sets.

other classes, and that *Indoors* and *Arts* are harder to differentiate amongst each other, as both use artificial illumination.

To show that the separability is not due to different patterns of luminance/color amongst classes, we have devised an experiment to measure the expected value of the Ueli descriptor conditioned to the luminance of the region.

We use a local sliding window to calculate local luminance and Ueli, and compute them on all possible 32×32 -pixel windows in 300 examples of each class to estimate the $E[U_T | Lum, Class]$, the conditional expectation of Ueli for given luminance and class.

We approximate the continuous function using histograms of expected values of U_T for each class H_i^E , $i \in \{Outdoor, Indoor\}$

$$H_i^E \leftarrow E[Ueli | L \in 1 \dots 255], \quad (4)$$

where $E[\cdot]$ is the statistical expectation, and L is the luminance such that $L = (0.3 * R) + (0.59 * G) + (0.11 * B)$.

The upper plot of Figure 12 displays the Ueli conditional expectation for unmodified *Outdoors* and *Indoors* classes. There is a consistent difference between classes, showing that the separability of the LSB statistical descriptors is not due to different class patterns of luminance.

We also observe the effect of the limited dynamic range on the statistical descriptor. Luminance components that are too small are squished to zero, while color components that should be very high are clamped to the maximum (255 in the 8-bit case). In these extreme cases, there is no randomness, and the Ueli value goes down to zero. As we calculate the expected values in sliding windows, the decrease along the borders of the dynamic range demonstrate the decrease of the randomness as more elements of the window are clamped to an extreme value.

Further validation with high dynamic range cameras, and with illumination controlled experiments (both out of the

scope of this work), would help us have a better understanding of the technique.

5. Conclusion and remarks

We have introduced a methodology that captures the changing dynamics of the statistical artifacts inserted during a perturbation process in each of the broad-image classes of our interest.

We have applied and validated the Progressive Randomization framework to the *Broad-Images Categorization* problem. Unlike other approaches of the literature, our method works in a multiclass scenario and differentiates several broad-classes of images using a single descriptor that has a relatively small dimension.

The most important features in the PR descriptor are its low dimensionality and its unified approach for different applications (e.g., the class of an image, the class of an object in a restricted domain) even with different cameras and illumination.

We have showed, experimentally, that different classes do have a distinct behavior under PR. Our conjecture is that the interaction of different light spectrum with the camera sensors induces different patterns in the LSB field. PR does not consider semantical information about scenes.

With enough training examples, the PR descriptor can be used as a complete self-contained classification procedure. However, huge training sets are not always available. In this case, with few training examples, PR descriptor presents two interesting properties that indicate that it can be combined with other image descriptors such as those described earlier in this paper. First, it generalizes from the priors even for small training sets. Second, the accuracy increases with the number of training examples.

Here, we illustrate the use of PR descriptor using LSB statistics. However, we can use other image channels be-




Figure 12. Dynamic ranges of the conditional Ueli descriptor given the luminance variation. Original image set, and the same image set with perturbations of 10%, 25%, and 50%, respectively. The plots are in different scales.

sides the LSB. Also, we can select image regions rich in details and analyze how they are affected using PR descriptor. Furthermore, we can use other statistical descriptors besides χ^2 and U_T such as kurtosis and skewness.

Future work include the application of PR to detect art-forgery, and to perform analysis and recognition of the type of sensor that took a picture.

6. Acknowledgments

The authors thank the financial support of FAPESP (Proc. 05/58103-3), and CNPq (Proc. 301278/2004).

References

- [1] E. L. Allwein, R. E. Schapire, and Y. Singer. Reducing multiclass to binary: a unifying approach for margin classifiers. *JMLR*, 1:113–141, 2001.
- [2] A. Bosch, A. Zisserman, and X. Munoz. Scene classification via pLSA. In *ECCV*, 2006.
- [3] F. Cutzu, R. Hammoud, and A. Leykin. Distinguishing paintings from photographs. *CVIU*, 100:249–273, 2005.
- [4] L. F. Fei, R. Fergus, and P. Perona. One-shot learning of object categories. *TPAMI*, 28(4):594–611, 2006.
- [5] D. Freedman, R. Pisani, and R. Purves. *Statistics*. George J. McLeod Limited, 1978.
- [6] J. Fridrich, M. Goljan, and R. Du. Reliable detection of LSB steganography in color and grayscale images. In *ACM Multimedia and Security*, pages 27–30, 2001.
- [7] J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*. Springer Verlag, 2001.
- [8] K. Grauman and T. Darrell. Efficient Image Matching with Distributions of Local Invariant Features. In *CVPR*, pages 627–634, 2005.
- [9] G. Heidemann. Unsupervised image categorization. *Image and Vision Computing*, 23:861–876, 2005.
- [10] J. Sivic and B. Russell and A. Efros and A. Zisserman and W. Freeman. Discovering objects and their location in images. In *ICCV*, pages 370–377, 2005.
- [11] J. Luo and A. Savakis. Indoor vs. outdoor classification of consumer photographs using low-level and semantic features. In *ICIP*, pages 745–748, 2001.
- [12] S. Lyu and H. Farid. How realistic is photorealistic? *IEEE Trans. on Signal Proc.*, 53:845–850, 2005.
- [13] M. Marszałek and C. Schmid. Spatial Weighting for Bag-of-Features. In *CVPR*, pages 2118–2125, 2006.
- [14] U. Maurer. A universal statistical test for random bit generators. *Journal of Cryptology*, 5:89–105, 1992.
- [15] A. Oliva and A. B. Torralba. Modeling the shape of the scene: A holistic representation of the spatial envelope. *IJCV*, 42(3):145–175, 2001.
- [16] A. Payne and S. Singh. Indoor vs. outdoor scene classification in digital photographs. *Pattern Recognition*, 38(10):1533–1545, 2005.
- [17] A. Rocha and S. Goldenstein. Progressive Randomization for Steganalysis. In *8th IEEE Intl. MMSP*, 2006.
- [18] N. Serrano, A. Savakis, and J. Luo. A computationally efficient approach to indoor/outdoor scene classification. In *16th ICPR*, pages 146–149, 2002.
- [19] A. Vailaya, A. Jain, and H. J. Zhang. On image classification: city images vs. landscapes. *Pattern Recognition*, 31:1921–1935, 1998.
- [20] J. Vogel and B. Schiele. A semantic typicality measure for natural scene categorization. In *DAGM Annual Pattern Recognition Symposium*, 2004.
- [21] P. Wayner. *Disappearing cryptography*. Morgan Kaufmann Publishers, 2002.
- [22] A. Westfeld and A. Pfitzmann. Attacks on steganographic systems. In *3rd Intl. Workshop on Information Hiding*, pages 61–76, 1999.